

Irrelevant Influences

KATIA VAVOVA

Mount Holyoke College

We often hear such casual accusations: you just believe that because you are a liberal, a Christian, an American, a woman... When such charges are made they are meant to sting—not just emotionally, but epistemically. But should they? It can be disturbing to learn that one's beliefs reflect the influence of such irrelevant factors. The pervasiveness of such influence has led some to worry that we are not justified in many of our beliefs. That same pervasiveness has led others to doubt whether there is any worry here at all. I argue that evidence of irrelevant belief influence is sometimes, but not always, undermining. My proposal picks out ordinary, non-skeptical cases in which we get evidence of error. It says that, in those cases, evidence of irrelevant influence is epistemically significant. It shows how respecting evidence of error is compatible with the epistemic lives we see ourselves living. We are fallible creatures, yes, but we are also capable and intelligent ones. We can recognize and correct for our own error so as to improve our imperfect, yet nevertheless robust, epistemic lives.

1. The challenge

In 1961, G. A. Cohen chose Oxford over Harvard for graduate school. He later realized that philosophers of his generation who studied at Oxford tend to accept the analytic/synthetic distinction, while those who studied at Harvard tend to reject it—despite being privy to all the same arguments. This lead Cohen to worry that “in some sense of ‘because’, and in some sense of ‘Oxford’”, he accepts the analytic/synthetic distinction because he studied at Oxford. “And that”, he thinks, “is disturbing”, because the fact that he studied at Oxford is “no reason to think that the distinction is sound”.¹

Cohen's position is familiar. The fact that you were raised in this community rather than that one is neither here nor there when it comes to what you ought to believe about God, morality, or presidential candidates. Yet factors like upbringing inevitably guide our convictions on these and other, less charged, topics. The effect is not always straightforward—perhaps you wouldn't be so liberal if you hadn't been raised in a liberal household, or perhaps you wouldn't be such a staunch atheist if your parents hadn't been so profoundly religious—but it is disturbing either way.

Irrelevant factors are factors that don't bear on the truth of what we believe. We rightly omit such factors when giving reasons. It's tempting to think that we should believe what we do because of evidence and arguments—not because of where we were born, how we were raised, or what school we happened to attend. If that is right,

¹ Cohen (2000: 18).

however, and if such influences really are pervasive, then we are irrational in much of what we believe. But that seems absurd—it doesn't seem irrational, in general, for thoughtful, careful people to have thoughtful, careful beliefs about even the most controversial matters. As Cohen notes, “we do not normally consider beliefs of [this] sort as instances of our (perhaps even commonplace) irrationality”.²

Learning that a belief reflects the influence of an irrelevant factor is troubling. But such influence is pervasive. We are thus left with a puzzle. It can't be that any belief influenced by irrelevant factors is irrational or we'd be left with mass skepticism. Yet it also can't be that no such beliefs are in trouble—some irrelevant influences seem truly pernicious. To solve this, we need to understand what it is about evidence of irrelevant influences that is so threatening when it is. We need a way to distinguish between problematic and innocuous cases. Without that, we face a serious problem about what to believe—specifically, about whether and how to go on believing what we were raised to believe, once we realize that we were raised to believe it. We need to know when, how, and why learning that a belief reflects the influence of irrelevant factors should affect our rational confidence in that belief.

While the importance of this question has been noted, surprisingly little has been said to directly address it.³ Much of what has been said is dismissive or reductive.⁴ Even Cohen finds the puzzle too puzzling, suggesting that we simply set aside such “morose meanderings”.⁵ The primary goal of this paper is to remedy this neglect. My proposed remedy has two parts. The first is a defense of the problem. I argue that evidence of irrelevant influence poses a unique and substantial epistemic challenge that does not collapse into other familiar challenges like disagreement or skepticism. The second is a proposal: a principled way of thinking about irrelevant influences that distinguishes problematic from innocuous cases. I argue that evidence of irrelevant influence is sometimes evidence of error. Rationality demands we not ignore it. The aim of this paper is thus to shed some light on the nature and import of such evidence: what it is and what it is to rationally accommodate it.

A final clarification before I proceed. Although my proposal provides a framework for thinking about this problem, it doesn't make it easy to know, in any given case, whether a belief influence is innocuous or problematic. Nor does it tell us exactly what to do if we discover a problematic one. Rather, my primary goal is to explain the difference between epistemically significant and insignificant belief influences, and thereby resolve some of our philosophical puzzlement. Determining what to believe in a given scenario will still require substantive thinking and we may remain, in the end, substantively puzzled. But that's fine. The end of philosophical puzzlement is not the end of thinking.

² Cohen (2000: 13).

³ For work that discusses these issues more or less directly see Anur & Scott-Kakures (forthcoming), De Cruz (ms), DiPaolo and Simpson (2015), Dworkin (1996, 2011), Elga (ms), Mogensen (2015), Rosen (2001), Schechter (ms), Schoenfield (2012), Sher (2001), and White (2010). For similar themes in philosophy of religion see Garber (2007, 2009) and Plantinga (2000). On related Darwinian doubts about morality see Vavova (2015)'s overview. For the continental side see Leiter (2004)'s discussion of Marx, Nietzsche, and Freud.

⁴ See especially Dworkin (1996, 2011), Mogensen (2015), and White (2010), whom I'll discuss shortly.

⁵ Cohen (2000: 19).

2. The influences

A belief influence is irrelevant relative to a belief and a person. In other words:

An **irrelevant influence** for me with respect to my belief that p is one that (a) has influenced my belief that p and (b) does not bear on the truth of p .

This relativity makes sense: that Elise was born in San Francisco may have influenced her beliefs about gay rights, but it is strictly irrelevant to the truth of those beliefs. That same fact about her birthplace, however, *is* clearly relevant to *my* belief about where Elise was born. So the same factor can be irrelevant with respect to one belief or for one person and relevant for another belief or another person. This shows that there isn't one particular sort of factor, be it genetic, social, historical, or whatever, that will always count as irrelevant.

There is, furthermore, no single way in which irrelevant factors might influence our beliefs. Sometimes, the influence might be causal, so that some factor F caused me to believe p . Other times, the influence is better described counterfactually: even if F didn't cause me to believe that p , it can still be troubling to realize that if it hadn't been for F , I wouldn't have believed that p . (The "you only believe that. . ." formulation suggests this.) Yet, in other cases, this counterfactual formulation fails: it seems true that we believe p because of F , in the counterfactual sense of 'because', and yet there is no counterfactual scenario to speak of: e.g., if it hadn't been for F , I wouldn't even have existed!⁶ Some have described the phenomenon as one about contingency or accidentally.⁷ But this doesn't capture all instances of problematic influence. There are cases of irrelevant influence in which it was no accident that I believed what I do—maybe it was even necessary. Yet my belief still seems in trouble.⁸

Because irrelevant influences admit of such variety, there isn't much we can add to the minimal definition above without excluding relevant cases. We can, at best, give examples of problematic cases. Such cases include, at least, intellectual, social, genetic, or evolutionary influence.⁹ Although the mechanisms at work are different in each case, they have something in common: they are all cases in which some factor, which itself doesn't bear on the truth of p , influenced our belief that p . The minimal definition allows us to get a handle on the phenomenon we are interested in, in all its diversity.

It also helps us separate two distinct questions about irrelevant influences:

1. Which belief influences are *irrelevant*?
2. Which belief influences are *epistemically problematic*?

Our answer to (1) shouldn't automatically settle the answer to (2). If, for example, we gave an answer to (1) that characterized irrelevant influences as epistemically problematic by their very nature, we would beg the question against theorists who argue that evidence

⁶ Cases of evolutionary influence on belief are often described in this way. See Street (2006) and Vavova (2014a).

⁷ Cf. Mogensen (2015).

⁸ DiPaolo and Simpson (2015)'s characterization in terms of indoctrination is similarly narrow.

⁹ De Cruz (ms) also cites cases of confirmation bias, wishful thinking, and self-selection.

of irrelevant influence doesn't always call for belief revision.¹⁰ The minimal definition above doesn't do this. It answers only the first question, and thereby respects an important fact about irrelevant influences: they're not *all* bad.¹¹ The above definition is thus just right for our purposes: it captures the relevant phenomenon in a theoretically neutral way.¹²

3. The challenge is challenging

I think this is a real problem, but you may be tempted to dismiss it. I'll therefore start by responding to the worry that either that there is no problem here, or that there is, but it is not a big deal. I'll consider three suggestions, all of which aim to show that worries about irrelevant influences are either misplaced or overblown. The first focuses on the fact that irrelevant influences are, after all, *irrelevant*. So, it concludes, discovering them shouldn't affect what we believe. The second move argues that rationality is permissive—permissive enough, at least, so that evidence of irrelevant factors is not undermining. The third aims to show that the problem of irrelevant influences reduces to that of disagreement. In a later section I'll consider the hypothesis that awareness of irrelevant influences merely makes salient more general skeptical worries, but does no undermining work on its own. According to the last two suggestions, there may be a real challenge here, but it is either uninteresting or unoriginal. According to the first two, there is no challenge at all.

My responses in this section will be quick and a bit promissory. They will, however, warm us up to the idea that this is a real problem—one that connects to other important epistemological problems. I hope, by the end of the paper, to have made good on my promises by showing that all four attempts to dismiss the challenge fail. For now, return to Cohen.

3.1. Evidence of irrelevant influence is not irrelevant

Cohen realizes that if he had gone to Harvard, he would deny what he now believes despite being privy to the same considerations. This is unnerving, but does it, on pain of irrationality, require Cohen to revise his confidence in or abandon his belief that the analytic/synthetic distinction is sound? If it does, then how are we ever justified in believing anything? (Every one of our beliefs reflects the influence of *some* irrelevant factor.) If it does not, then why is it so unnerving? (What are we reacting to and are we overreacting?) This is Cohen's problem. It feels like a real one, but what, if anything, gives his realization its skeptical bite?

¹⁰ Other available characterizations of 'irrelevant influences' fall short in a few ways. Some explicitly don't define or circumscribe the phenomenon, but rather rely on examples to give the reader an intuitive grasp (e.g., Schoenfield (2012) and White (2010)). Others are either insufficiently neutral or insufficiently inclusive (e.g., Mogensen (2015), diPaolo and Simpson (2015)).

¹¹ Another virtue of separating these questions (and answers) is that just as not all irrelevant influences are epistemically problematic, not all cases of epistemically problematic evidence are cases of irrelevant influences. Irrelevant influences, when they are problematic, are instances of a more general phenomenon, as I will argue shortly. Ideally, then, answers to (2) should apply more broadly, not just to cases of irrelevant influences.

¹² Thanks to an anonymous reviewer for urging me to make the points in this section.

Perhaps nothing does. Perhaps Cohen's realization has no skeptical bite. True, Cohen has evidence that he would have believed otherwise had he attended Harvard. But how could that possibly bear on what he should believe? Either Cohen has good reasons for his belief, or he doesn't. If he doesn't, he should abandon it. If he does, then the fact that he would have believed otherwise seems irrelevant. The good reasons he had for his belief remain. Finding out that you believe something because you were raised to may be surprising, but how could it be evidence against your belief?

If you have indeed formed your belief on good evidence, it can be hard to see why learning that it reflects the influence of some irrelevant factor should have any effect on what you believe. Why should it make you turn your back on all the perfectly good evidence that originally led you to believe as you do? Why can't you maintain your belief and, as Ronald Dworkin suggests, simply "count it as a piece of luck" that your upbringing inclined you toward believing what your evidence supports?¹³

On this line of thought, Cohen needn't worry. He can just thank his lucky stars. The thought is that facts about the nature and quality of our evidence, our upbringing, and our manner of belief formation are wholly irrelevant to what we ought to believe about unrelated matters. We should never be moved by anything other than our original, first-order, evidence. Higher-order evidence, evidence about the quality of that first-order evidence or, relatedly, evidence of error, is never epistemically significant. It shouldn't affect what we believe.¹⁴

But evidence of error is still evidence. And it sometimes gives us reason to revise. Learning that you were brainwashed to believe something casts doubt on your belief. As does discovering that your favorite biology textbook is sorely outdated. Likewise, the discovery that the world's expert paleontologists disagree with you about the size of the average Barosaurus (big and heavy) should lead you to revise your belief that it is small and lithe. You ought to revise in this case, even if, by some fluke, you, the layperson, are right and they, the experts, are wrong. This is so even though sociological facts about opinions and publication dates are, strictly speaking, irrelevant to the truth of what you believe.

You might agree about the significance of evidence of our error while still disagreeing about these particular cases. That doesn't matter here. The point is about the Dworkonian position. It entails that I can ignore even the most damning evidence of error—so long as that evidence is evidence *about* my evidence rather than evidence pertaining directly to my first-order belief. But that is absurd.¹⁵ Learning that I was hypnotized to believe that Portland is the capital of Maine should cast doubt on that belief just as reading an outdated atlas that says it's Augusta should. It doesn't matter that the latter evidence is direct, first-order evidence while the former isn't. A position like this, which denies the epistemic significance of evidence of our error, cannot accommodate these considerations. So it must be wrong.

¹³ Dworkin (1996: 125).

¹⁴ Sher (2001: 70) summarizes but doesn't endorse this position. Dworkin (1996, 2011) endorses it. Kelly (2005) takes a similar position about disagreement. Cf. Weatherson (ms). For the opposite view about disagreement, see Elga (ms, 2007) and Christensen (2010, 2011).

¹⁵ Dworkin embraces these consequences (see especially his (2011: 77)). Kelly (2010) backs off from his earlier (2005) Dworkin-like views.

3.2. Going permissive doesn't help

Permissivism is the view that sometimes there is more than one rational response to a given body of evidence. The quick argument from permissivism to the epistemic insignificance of irrelevant influences goes like this: if both belief and disbelief in p are rational on my evidence, then why should I care that I could have believed otherwise? After all, my belief is rational (and would be either way).¹⁶

This is a way out of our puzzle, but why should we accept permissivism? There are at least two motivations. The first is anti-skeptical. If permissivism is true, then there are fewer threats on my beliefs: evidence of irrelevant influences or disagreement isn't evidence of irrationality. You and I may be equally thoughtful, intelligent people who disagree about the permissibility of abortion. Perhaps we only disagree because we were raised in politically different households. None of this need worry us, on this line of thought, if we're both equally rational. And, permissivism allows us to be—even if we share all the same evidence but come to different conclusions.

The second motivation for permissivism is intuitive. Gideon Rosen thinks it “should be obvious that reasonable people can disagree, even when confronted with a single body of evidence”.¹⁷ Rosen argues that when a jury is divided on a difficult case or when scientists disagree about what killed the dinosaurs, none of the disputants need be irrational. Similarly, do we want to say of Lewis and Stalnaker's disagreement about possible worlds that one of them is *irrational*? What about Quine, Strawson and their respective graduate students? Is it plausible that most of them are irrational?¹⁸ If not, then maybe Cohen has a way out.

Disagreements about tricky matters like these motivate permissivism. But even if permissivism is true, it doesn't follow that we needn't worry about irrelevant influences cases. Here's why. The quick argument above presupposes that the only thing evidence of irrelevant influence could be evidence of is irrationality. That's not obvious (it could, for example, be evidence of falsehood). Even if this is right, however, it doesn't follow that we can never get such evidence in permissive cases—not, at least, on plausible versions of permissivism. To see this, we need some background about permissivism.

On permissive views, the rationality relation has three prongs: the evidence, the belief, and the agent. Different versions of permissivism pick out different features of the agent as relevant. Options include the agent's prior probabilities, her fundamental inductive methods, or her epistemic standards.¹⁹ An agent's epistemic standards are sometimes described as the agent's way of taking evidence into account. Suppose you favor explanatory power above all theoretical virtues, while I favor simplicity above explanatory power. We might, then have different epistemic standards and the same evidence

¹⁶ This is too quick, but suffices for my purposes. See Schoenfield (2012) for a more careful version of this argument.

¹⁷ Rosen (2001: 4).

¹⁸ Do not be distracted by the focus on disagreement. The question here is about whether two people who evaluate the same bit of evidence can rationally come to different conclusions—it is not about what these people should do when they *learn* that they've come to different conclusions. The latter question is about the epistemic significance of disagreement; the former is about the truth of epistemic permissivism.

¹⁹ Subjective Bayesianism is formulated in terms of priors (Meacham (2014). Schoenfield (2012) and White (2005) talk in terms of epistemic standards. 'Inductive method' comes from Lewis (1971). 'Epistemic standards' is the least clear term. I'm not sure it is picking out something other than inductive methods or prior probabilities of a sort.

could lead us to accept different theories. If both our standards are rational, then it will be rational for us to accept different theories on the same evidence.

Here's why such permissivism doesn't preclude the possibility of evidence of irrationality. The rationality of what I believe is a function of my epistemic standards, yes, but I might violate those standards. Evidence that I am violating those standards is evidence of irrationality, and evidence of irrelevant influence can be such evidence. For example, I have consistency as an epistemic standard, but could be pushed by bias, poor memory, or wishful thinking into accepting inconsistent beliefs. If I then learned that I was flawed in some such way, that would give me reason to worry about my belief even if permissivism holds. This, incidentally, is how Miriam Schoenfield diagnoses Cohen's situation. Though it is a permissive case, and though Schoenfield thinks that permissive cases are, in general, immune to worries about irrelevant influences, she thinks that this case isn't. Evidence of irrelevant influence may give Cohen reason to think that he has reasoned irrationally, "by reasoning in ways that are inconsistent with [his] other beliefs and standards".²⁰

To sum up, it isn't clear that permissivism eliminates the worry that our epistemic standards are irrational. Only the most extreme form of permissivism, on which anything goes as far as rationality is concerned, could do this. Such a view eliminates worries about irrelevant influence because it altogether eliminates the possibility of evidence of error. Nothing puts pressure on rational agents, on this view. This crude a view, however, isn't what permissivists typically have in mind.²¹ More moderate permissivist views are more plausible the further they are from this extreme.²² But none of these more plausible views eliminate the worry about irrelevant influences. At best, they cut down the number of ways in which we might get evidence of our error. So long as they (rightly) don't do away with such evidence entirely—so long as evidence of irrelevant influence can be evidence of some sort of irrationality or error, like violating one's epistemic standards—it remains a threat even in permissive cases.

3.3. *Not just disagreement*

Cohen's concern is initially prompted by discovering a correlation: Oxford philosophers tend to accept the analytic/synthetic distinction while Harvard philosophers tend to reject it. It is this evidence—evidence of disagreement that first worries Cohen. Perhaps, then, Cohen's problem isn't that he would have believed otherwise, but rather that other intelligent people do.

Roger White argues that this is exactly right—that whatever skeptical force considerations about irrelevant belief influences have, they only do so by piggybacking on other factors with skeptical force, such as disagreement.²³ But aren't there disagreement-free cases of irrelevant influences? Anticipating this, White considers a variant of Cohen's situation in which there is no disagreement, yet Cohen is rightly worried about his belief.²⁴ In this scenario, there are no Harvard graduate students. However, Cohen has good

²⁰ Schoenfield (2012: 17).

²¹ Cf. Christensen (2014: 6).

²² Compare Kelly (2013), Meacham (2014), and Schoenfield (2012).

²³ White (2010); cf. Mogensen (2015).

²⁴ White (2010: 607).

reason to believe that if Harvard had graduate students, they would reject the analytic/synthetic distinction. White grants that this undermines Cohen's belief in the distinction. He argues, however, that plausible ways of explaining this residual worry will still collapse it into a worry about disagreement. To show this he considers a few ways we might explain Cohen's worry.

Perhaps, e.g., Cohen's worry remains because he is familiar with the ways in which graduate student opinion can be swayed by professor's opinion in less than rational ways. Or perhaps he realizes that intelligent, informed, and articulate philosophers present strong and compelling cases for their views—even when those views are false. So if Harvard professors reject the analytic/synthetic distinction, then their merely possible Harvard students are likely to also reject that distinction. (The same may be said of Oxford, of course, and this is the problem.) True, these Harvard students don't exist, but that is irrelevant.²⁵ The worry is that *if* we were to send smart, thoughtful graduate students through Harvard, they would come out rejecting the analytic/synthetic distinction. And that suggests that the Oxford students might be mistaken to accept it. White thinks that if explanations like these are the best we can do, then "[W]e have really just come back to the issue of disagreement".²⁶ So, he concludes, whatever epistemic significance of evidence of irrelevant influence has, it derives solely from considerations about disagreement.

But this is not the best we can do and disagreement is not our culprit. Perhaps every case of irrelevant influence is also a case of disagreement in this weak sense: it is a case of *possible* disagreement (with, at least, the possible you that wasn't so influenced). But this doesn't establish much. First, as White himself argues, it isn't clear that merely possible disagreement should worry us. Someone who wants to defend White's disagreement hypothesis needs to explain (a) why we should worry about these merely possible disagreements at all, and (b) how this doesn't collapse into a general skeptical worry, given that there is always *possible* disagreement.²⁷ Second, even if wherever there are irrelevant influences there is (at least possible) disagreement, it doesn't follow that the worry comes *from* or *is about* the (perhaps imagined) disagreement. At best, this establishes a correlation—not causation.

There is, furthermore, a better way to understand the relationship between disagreement and irrelevant influences. Even if every case of irrelevant influence is also a case of (at least possible) disagreement, not every case of disagreement is a case of irrelevant influence. And these two sorts of cases often suggest different sorts of error. Some disagreements, like simple arithmetical ones, are best explained in terms of simple error (one of us miscalculated, as fallible creatures are wont to do). Other disagreements are best explained in terms of systematic and pervasive error of the sort that evidence of irrelevant influence typically points to (bias, inculcation, etc.). This suggests that evidence

²⁵ This disagreement is what Mogensen (2015) calls "arbitrarily absent". Unlike merely possible disagreement, arbitrarily absent disagreement can still be evidence of error. I agree with Mogensen on this point, but not on his defense of what he calls the 'Disagreement Hypothesis'—White's view that the problem of irrelevant influences is just a problem about disagreement.

²⁶ White (2010: 608).

²⁷ Mogensen (2015) does a nice job of this. However, even if he is right that we should worry about some possible disagreements, like those in problematic irrelevant influence cases, it doesn't follow that the problem of irrelevant influences is simply a problem about disagreement for the reason stated above: we have at best correlation, not causation.

of irrelevant influence and evidence of disagreement are species of the same genus: evidence of error.

Evidence of error comes in all shapes and sizes and its species are closely related: sometimes evidence of disagreement will be what alerts us to the presence of irrelevant influence—other times it'll be the other way around. But there are other ways we might learn about such influence. We might read the latest studies in cognitive psychology, which show just how many distorting factors affect our beliefs. We could find out that we have been systematically brainwashed by an evil dictator, or, more realistically, that we are motivated by advertisements. Alternatively, we might discover that evolutionary forces inclined us toward certain beliefs not because they are true, but because believing them promotes fitness. This discovery might give us evidence of detrimental belief influence. Notice, crucially, that it might do so even if we do not think that there is any sense in which we would have believed otherwise. The closest possible world in which evolutionary forces work differently may be one in which we have no recognizable counterparts that could disagree with us (perhaps because evolution quit at slugs). This doesn't dissolve the worry that evolution may have had an adverse affect on our thinking.²⁸ Likewise, for the fact that no one disagrees with the dictator and the fact that all of us are biased in some of the same ways. The lack of dissenting opinion—possible or actual—does nothing to dissolve the worry that our cognitive faculties can't be trusted.

The worry about irrelevant influences is thus not that someone with your evidence believes otherwise. Nor is it exactly that an alternate *you* would have believed otherwise. It is that your actual beliefs may not be reliably formed. This isn't a worry about disagreement or skepticism. It's a worry about irrationality or error.

4. Diagnosis

Consider a case of irrelevant influence with some nice simplifying features: (1) it concerns what is clearly a substantive matter, (2) it is clear what the relevant evidence is, (3) not all responses to that evidence are equally rational, and (4) it is clear what the irrelevant influence is.

Primed. You participate in an experiment. You are shown a short video and then given some unrelated data on pet therapy for chronically ill elderly patients. This is a topic about which you know nothing. You find the data compelling and leave the lab newly convinced that pet therapy is an effective way to relieve patient suffering. "Grandma hates animals," you think to yourself, "but the data is just overwhelming. I'd better stop by the pet store." You come home, pooch in pocket, to the following email from the experimenters:

Thank you again for participating in our study. The goal was to determine the effectiveness of visual priming. Here's what we did. We split subjects into two groups. We used subliminal visual cues in the introductory video to prime half toward a particular interpretation of the data. The other group was not primed; their video contained no subliminal messages. We then presented both groups with the same data. We found a striking correlation: all the primed subjects believed as they were primed to and thought that the data provided compelling grounds for their conclusions. We

²⁸ Joyce (2006) and Street (2006) defend versions of this evolutionary argument. For an overview see Vavova (2015).

cannot disclose to which group you, Sir or Madam, were assigned. Nevertheless, we must warn you that your confidence in your belief may well be overrated.

Suppose that this is your situation. You believe that Grandma would benefit from a puppy. Perhaps she would. Perhaps you were lucky to be in the unprimed group and did, in fact, respond to the data rationally. However, you have some reason to think that you were not so lucky: half of the subjects were primed to misread the data and there is a good chance that you were one of them.

This should worry you. You can't dismiss the experimenter's note with, "Well, pet therapy is effective, and I do believe that it is. I remember that compelling data they showed me. How lucky that I saw the evidence for what it was—how lucky that I wasn't primed!"²⁹ From the subject's position, this is a terrible way to dismiss evidence of irrelevant influence. But such a response is not always inappropriate. For example, the experimenter can legitimately dismiss the possibility of error with exactly the same words: "How lucky that I wasn't primed!"

Why is this response appropriate for the experimenter, but not for the experimented upon? After all, both subject and experimenter know that priming is akin to brainwashing or biasing: it makes one respond poorly to evidence. And it is true of both subject and experimenter that they *might* be primed. Here is the crucial difference: the subject has good reason to think that she has been primed—she has evidence that there is a 50 percent chance of it. The experimenter has no such reason. The experimenter knows that she *might* be primed, of course. She knows this in the same way in which she knows that she might be a brain in a vat, adopted, part Canadian, and so on. But in learning about the experiment, the subject learns that it is not only possible, but also *probable* that she was primed. It is appropriate for her to be more worried about her epistemic situation after she acquires this new information because it makes the possibility of her error more probable. The experimenter has acquired no such information, so she needn't worry.

To see the crucial difference compare the following scenarios.

Possible BIV. We are chatting over a cup of coffee. Out of the blue I say, "Wouldn't it be odd if we were really just brains in vats stimulated to have an experience as of having this conversation? We might be, you know".

Probable BIV. We are chatting over a cup of coffee. Suddenly, I confess: "Yesterday, I flipped a coin to determine whether to test my new envatting technology out on you. Heads, and I would seamlessly envat you and stimulate your brain to have an experience as of us having this conversation. Tails, I would just take you out for coffee. I can't tell you how the coin landed, but out of curiosity, how do you feel?"

Set aside how weird these cases are and notice: whatever doubt raising the skeptical possibility casts on your beliefs in Possible BIV, raising that same possibility casts a distinct sort of doubt in Probable BIV. This is because in the first scenario I raise the mere possibility of error. In the second I give you evidence that this possibility is probable: there is a 50 percent chance that you have been envatted. The second of these scenarios is more

²⁹ Not everyone thinks such epistemic luck is problematic. For example, Dworkin (1996), whom I discussed in section 3, thinks such a response is just fine. See also Plantinga (2000)'s related discussion of the noetic effects of sin.

worrisome than the first because it provides **new information**—information that there is a good chance you are in a bad epistemic situation.

The subject's position in Primed is similar. Before she learns about the experiment, she knows that she might be primed, that it is *possible* she has been primed (again, just as it is possible that she is adopted, part Canadian, etc.). Once she gets good reason to think that she was primed, the subject should be more worried about her epistemic situation. This is because what was a mere possibility of error has now been made more probable—roughly 50 percent more probable.³⁰ This is thus a crucial feature of the subject's situation in Primed: the evidence of belief influence she acquires gives her good evidence that she is in a bad epistemic situation. This shows how evidence of irrelevant influence does something more than merely make salient a general skeptical challenge. This also suggests a diagnosis: evidence of irrelevant influence is worrying when it gives us reason to think we are in a bad epistemic situation.

There are, of course, many ways in which our epistemic situation might be bad or we might be, as I've put it above, in error: we might be irrational, inaccurate, unreliable, incoherent, violating our epistemic standards, etc. These come apart—I might, e.g., be perfectly rational but totally inaccurate—and evidence of irrelevant influence could be evidence of any of them. For lack of a better term to encompass them all, I'll talk here of being *mistaken*. Being mistaken, in this technical sense, is not an all or nothing matter and it needn't apply only to beliefs. Evidence that I am mistaken could be evidence that I have a false belief, but it could also be evidence that I am irrational, overconfident, unreliable, incoherent, and the like.

5. Proposal

From the foregoing, we get an initial suggestion for how to deal with evidence of irrelevant belief influence:

Evaluate your epistemic situation. To the extent that this evaluation suggests that you are mistaken, adjust the influenced beliefs appropriately.

Notice that not just any evaluation will do. It is important that you not stack the deck in your favor—that you properly evaluate the likelihood of your own error. The right way of evaluating your epistemic situation should thus block fishy responses like: “I might be primed, eh? Well, but pet therapy *is* effective and I believe that it is! How lucky!” Once you acquire evidence that you are unreliable about some matter, you cannot assure yourself of your reliability merely by consulting your opinion about that matter. This is more than question-begging, or dialectically inappropriate, it is irrational and ineffective. You must establish, on *independent* grounds, whether you have good reason to think you are

³⁰ I'm working with one number—50—in both cases for simplicity. If, however, the subject learned that there was a greater or lesser chance that she was primed, she would have to adjust her confidence accordingly. Intuitively, learning that there was a 70 percent chance that I was in the primed group should make me much more worried than learning that there was only a 30 percent chance. Even in the latter, more favorable case, however, I still get reason to worry. We can make better sense of all this in a graded-belief framework, but exploring that is beyond the scope of the paper. For a glimpse of the complications and how we might accommodate them in the context of a different, but related debate, see Vavova (2014b).

mistaken—whether the irrelevant factor has had a detrimental influence on your belief. We can thus extract the following general principle:

Good Independent Reason Principle (GIRP). To the extent that you have good independent reason to think that you are mistaken with respect to p , you must revise your confidence in p accordingly—insofar as you can.³¹

Some clarifications. First, what must your reason be independent of? The p that has been called into doubt and the arguments and evidence for it. This is because, recall, you must evaluate your epistemic situation in a way that doesn't stack the deck in your favor. This means you shouldn't assume exactly what has been called into doubt.

Second, how are you meant to revise? Nothing more specific than 'accordingly' can be said at this general level. Whether your confidence moves up, down, side to side, or stays the same depends on the details. This is because you cannot, in general, correct for an untoward influence without knowing how it has influenced you. There may even be cases in which you cannot correct at all because you cannot know how you've been influenced, such as when you get very general evidence of error. This is why GIRP reads, revise "accordingly—insofar as you can".

Third, how significantly must you revise? Again, nothing very general can be said. It depends on how good of a reason you have for thinking that this influence is a bad one, how bad of an influence you think it is, how much reason you have for thinking that it has influenced you, how much you think it has, and etc. You may have a very strong reason to become just a smidgen less confident—or you may have a very weak reason to become substantially less confident. All this imprecision is unproblematic. Compare: I may not know how much I ought to donate to charity because I either have a very weak reason to donate a lot, or a very strong reason to donate a little. Regardless, I know I need to donate. Figuring out how much is the hard work left over when we're done sorting out the substantive moral questions. Likewise, even if we answer the epistemological question about when evidence of irrelevant influences is problematic, we may still wonder how exactly we must revise in any particular case. That will depend on the details.

Fourth, GIRP is quite general. It applies whenever we get reason to think we are mistaken, whether that evidence comes from disagreement, irrelevant influences, or whatever. This reflects the connection between evidence of irrelevant influences and evidence of disagreement, and underscores the way in which they are species of the same genus.³²

Fifth, since the principle activates when you get evidence that you are mistaken, it is applicable on any theory that allows for the epistemic significance evidence of your own error. This, recall, includes all but the most extreme permissivism. Insofar as evidence of irrelevant influence can be evidence of error, then, even permissivists (most of them) can embrace this principle.

Finally, on this proposal, the problem of irrelevant influence is importantly distinct from the problem of skepticism. The proposed principle, GIRP, is formulated so as to

³¹ Similar principles are explicit in Christensen (2007, 2009, 2011), Elga (2007), and Vavova (2014a). Christensen and Elga's discussions are about disagreement; Vavova's is about evolutionary debunking arguments. More must be said about what counts as independent and how to set aside what is not independent. Also, it is not obvious how to characterize this 'setting aside' formally. So, there is work left to do. Fortunately, these difficulties don't infect our discussions here (cf. Schoenfield (2012: fn. 28)).

³² Cf. section 3.3.

pick out evidence of ordinary sorts of non-skeptical error. It isn't activated by skeptical possibilities. I'll say more about this shortly, after testing the principle on a few cases.

6. Testing the proposal

I have argued that evidence of irrelevant belief influence is undermining to the extent that your independent evaluation gives you good reason to think that the influence has been detrimental—that it has lead you to be mistaken about the relevant matter. I argued that this is exactly the situation of the primed subject. How does this diagnosis help with more complicated real-life cases? To test it out, consider a few variations on Cohen's situation.

6.1. When Cohen should worry

As Cohen himself points out, "...one thing that I must believe [...] if I stick to my belief in the analytic/synthetic distinction, is that I am *lucky* with respect to the view I have of this matter".³³ But it is not enough for Cohen to think that he got lucky since the analytic/synthetic distinction is sound and he believes it. If the foregoing is right, Cohen cannot rationally respond in this way. He must determine whether his independent evaluation gives him good reason to think that he is mistaken about this matter. To be independent, this evaluation cannot be based on either his acceptance of the analytic/synthetic distinction, or on the arguments that led him to accept the distinction. So, does Cohen have a good reason to think that he is not reliable about the analytic/synthetic distinction?

Suppose, first, that he does have such a reason. Cohen might, for example, know that Oxford and Harvard are *not* epistemically on a par—that one is significantly better than the other. Perhaps the administrations flipped a coin to determine which of two chefs to hire. Both are excellent cooks, but one of them is nutty and known to sprinkle anti-depressants on the food. (Chef means well, she just wants the academics to be happy.) Common side effects include sleepiness, anxiety, and holding false beliefs that cannot be changed by fact. Both chefs get hired—one at each institution. (Times are tough and good chefs are rare.) Cohen does not know where the chefs ended up, but he knows that one group is doing philosophy under the influence.

This case isn't as silly as it might seem. Holding false beliefs that cannot be changed by fact is a listed side effect of some anti-depressants.³⁴ Likewise for sleepiness and anxiety—though notice, these are common side effects of being a philosopher too. This is important for our case because it makes it harder for Cohen to know if he's been drugged or if he's merely a tired, nervous academic.

In this case, then Cohen is situated just as the subject in Primed: he knows that one situation (the unprimed or drugless one) is better. He also knows that there is a 50 percent chance he is in the worse situation. This should lead the possibly primed subject to lose confidence in what she believes. The analogous discovery should do the same for Cohen: his independent evaluation of his situation gives him good reason to worry.³⁵

³³ Cohen (2000: 18).

³⁴ See, e.g., <http://www.drugs.com/sfx/celexa-side-effects.html>.

³⁵ Cohen's situation here is actually worse than the Primed subject's because he doesn't have a sense of which beliefs might be the false ones—presumably, it isn't just his belief about the analytic/synthetic distinction that is in trouble. Furthermore, given that his beliefs might not be responding to fact, it's not clear how he could revise. This is not a problem for the proposal. It merely demonstrates that sometimes we get evidence of error that either we don't know what to do with or we cannot do anything with.

Cohen's actual situation is trickier, but also plausibly one in which he has good reason to think that he is mistaken. Notice how he describes it. Cohen chose Oxford over Harvard on what was essentially a whim: Europe sounded more exciting than Massachusetts. Otherwise, Cohen took Oxford and Harvard to be equally good. He had no reason to think that belief-distorting drugs were in greater supply at Harvard than at Oxford. As far as revealing philosophical truth and training philosophers, Cohen took the two schools to be on a par.

Importantly, his years at Oxford did not change his mind. He did not, even after acquiring it, take his Oxford education to provide him with or deprive him of some special skill or insight. He thought such skills and insights were as likely to be gained at Harvard.³⁶ All of this suggests that, independent of his beliefs about the analytic/synthetic distinction, Cohen took Oxford and Harvard to be on a par. Does this give him good reason to think that attending Oxford increased his likelihood of being mistaken?

It does if we suppose that there is a fact of the matter about the analytic/synthetic distinction. The correlation between belief about the distinction and attendance at Oxford or Harvard suggests something suspicious is going on. It suggests that what graduate school one attends affects what one believes about the analytic/synthetic distinction. It does so despite the fact that neither school is epistemically privileged in any relevant way. So there must be a detrimental influence at play. Given that Cohen's independent evaluation of the situation gives him good reason to think that the two schools are on a par, and given that there must be error somewhere, that same evaluation gives Cohen reason to think that the error could lie with him.

If what I have suggested so far is right, then, on this description of the case too, Cohen cannot maintain confidence in his belief. To do so, he would have to reason as follows: "I had all the reason in the world to think that I would have been as well placed to determine whether p if I'd gone to Harvard. But I must have been wrong about that. After all, the analytic/synthetic distinction is sound, and if I'd gone to Harvard, I would have denied that". But Cohen shouldn't reason like this. So he must revise.

6.2. *When Cohen needn't worry*

According to GIRP, if Cohen's independent evaluation fails to give him good reason to think that he is mistaken, then he needn't worry. Here is one version of Cohen's situation that meets this condition.

When Cohen worried that he only believed in the analytic/synthetic distinction because he studied at Oxford, he may have been worried that he was unduly swayed by his advisor's opinions. This hypothesis would explain both why graduate students at each institution all come out believing both the same as each other *and* the same as their mentors do. Students may be in danger of falling prey to a guru effect of the sort cultists succumb to, becoming irrational or downright loony when under the influence of their charismatic leaders. The hypothesis only helps Cohen if he has reason to think that Harvard students are more likely to be affected by the views of their mentors. I have no idea if he had such reason, but we could fill in the details as follows, so that he does.

If Quine were super charismatic and Ayer a total bore, then Cohen could reasonably conclude the Harvard students are more likely to fall prey to the guru effect. Similarly, if

³⁶ Cohen (2000: 17).

such personalities are more powerful at smaller, more concentrated philosophical communities, then Cohen could worry less about their effects at less centralized places like Oxford. I'm not suggesting this is actually how things were. My point is that if things had been thus, and Cohen had reason to suspect so, then again, he might be able to thank his lucky stars.

Such considerations are appropriately independent of Cohen's beliefs about the analytic/synthetic distinction so long as the same story cannot be told about Oxford—or, at least, if the biasing influence of graduate school is more pronounced at Harvard. Notice also that these considerations would still be appropriately independent even if they, in a way, depend on Cohen's having studied at Oxford. Perhaps there is a psychology class on the Oxford curriculum that teaches about guru effects, which Cohen would not have taken if he had gone to Harvard. His beliefs about guru effects depend on his having gone to Oxford, but are independent of the relevant bit of his Oxford background: the analytic/synthetic distinction bit. What is crucial is that Cohen's reason for taking his Oxford situation to be superior has nothing to do with the fact that he would have believed otherwise if he had gone to Harvard. It has nothing to do with either his belief that the analytic/synthetic distinction is sound, or the arguments and evidence on which that belief is based. If it did, Cohen would be guilty of that same question-begging that sticking to your guns in the Primed case seems to require. In this case, then, Cohen's independent evaluation of his situation fails to give him reason to think that he is mistaken. So he needn't revise.

7. Not just skepticism

The best cases of irrelevant influence—the ones in which we can most easily maintain our beliefs—are ones in which we not only lack good reason to think we are mistaken, but we have good reason to think we are right. This is more than what Cohen has above, however, so you might worry that he doesn't have enough. Does the mere absence of reason to think we're mistaken suffice for rationally dismissing the worry that we are mistaken? Or do we need better assurance than this?

That we need more is a natural and compelling thought. If it is right, then we were too quick to accept GIRP and should accept something like this instead:

No Independent Reason Principle (NIRP). To the extent that you *fail* to have good independent reason to think that you are *not mistaken* with respect to matters like *p*, you must revise your confidence in *p* accordingly—insofar as you can.³⁷

White thinks that this is exactly the sort of principle that underlies Cohen-type worries. It is natural to think, as he puts it, that we need “something *independent* of the reasoning that lead [us] to *p* on which to judge that [we have] reached the truth [about *p*]”.³⁸ White is suspicious of such a requirement because it is “familiar from traditional skeptical arguments”.³⁹

³⁷ Cf. Christensen (2011: 25–26), Street (2015), and Vavova (2014a).

³⁸ White (2010: 601).

³⁹ White (2010: 604). See also Elga (ms).

He is right to be suspicious. By demanding that we independently establish our own reliability, NIRP effectively requires us to have an independent reason to think we are not in a skeptical scenario. But such scenarios are constructed exactly so that there is nothing independent—nothing that isn't called into question. White concludes that if we are to avoid "a very general skepticism", it must sometimes be legitimate to "endorse our own reasoning from within"—in a non-independent way.⁴⁰

But this is too quick. We needn't legitimize such patently poor reasoning. We can accept a more narrow, targeted skepticism without giving up all that we believe. The trick is to accept a more modest independence requirement. White only considers NIRP-like principles, on which we cannot maintain our beliefs if we *lack* independent reason to think that we are *reliable*.⁴¹ GIRP is a plausible and motivated alternative. It only requires we revise if we *have* independent reason to think that we are *unreliable*.

This is the right sort of principle for our purposes. Evidence of irrelevant influence indicates an ordinary sort of error—the sort we succumb to because we are fallible. This is the sort of error to which a more modest independence principle like GIRP is sensitive. Furthermore, NIRP doesn't distinguish between possible and probable error. This is a strike against it. For the same reason, it cannot distinguish between the good and bad cases of irrelevant influences. GIRP has no such problems. It is activated only when we get apparent evidence of ordinary sorts of error. It says nothing about the extraordinary sort of error that the skeptic envisages. There, it remains silent.

A final variant of Cohen's situation shows an important way in which NIRP and GIRP rule differently. Suppose that Cohen's entire body of belief rests on his conviction that the analytic/synthetic distinction is sound. As I am imagining it, Cohen's beliefs form an upside down pyramid, with his belief in the analytic/synthetic distinction at the very bottom. This is Cohen's most fundamental belief. It forms the foundation for everything else he believes and itself rests on no further belief. Should evidence of irrelevant influence undermine Cohen's confidence in this scenario?

No. And this is what we get if we accept GIRP only. GIRP tells us to revise only if, and to the extent that, our independent evaluation gives us good reason to believe we are unreliable. To perform an appropriately independent evaluation in this case, Cohen would have to set aside his belief that the analytic/synthetic distinction is sound. But everything else he believes rests on that belief; so he would have to set aside everything he believes. This leaves Cohen unable to independently evaluate his situation. He cannot step outside all that he believes and determine whether it is justified. NIRP requires that he should, and then punishes him when he can't. GIRP doesn't do either.

Notice that it does not follow from GIRP that the possibility of global error is nothing to worry about. In fact, nothing about skepticism follows from GIRP. Since GIRP is only sensitive to local error, it remains silent on the question of what to do with global error. This is as it should be. The question here is how to evaluate a certain sort of evidence. But evidence is evidence only against a backdrop of beliefs we take for granted. If I can take nothing for granted, I cannot evaluate anything. When evidence of irrelevant influence is evidence of error, I can discover this by bracketing away the beliefs called into question and asking myself whether I have good reason to

⁴⁰ White (2010: 604).

⁴¹ Likewise with Elga (ms.).

doubt them. In some skeptical scenarios, it is unclear what evidence of error even looks like. Since the skeptic calls my entire body of belief into question, I have no independent ground from which to evaluate the supposed evidence that has been laid before me. I have no independent ground from which to evaluate anything. If I bracket everything, I cannot ask myself if I am likely to be wrong. I do not have the resources to do so.⁴²

This marks a crucial difference between ordinary and skeptical doubt: in the former case, but not in the latter, I am capable of independently evaluating the possibility that I might be wrong. Since the central question here is how to evaluate evidence that this possibility obtains, it is not *ad hoc* to restrict the proposal to the former, ordinary sort of error. It can be hard to know what to do with, or how to wrap our heads around, the possibility that we might be globally unreliable. But we can certainly make sense of the possibility that we are wrong about this or that thing. My proposal is one story about what we should do with the possibility of such local error—a story that distinguishes it in a principled way from the possibility of global error.

8. Conclusion

It can be disturbing to learn that one's beliefs reflect the influence of irrelevant factors. The pervasiveness of such influence has led some to worry that we are not justified in many of our beliefs. That same pervasiveness has led others to doubt whether there is any worry here at all. I argued that evidence of irrelevant belief influence can be undermining: we are fallible creatures, after all. I also argued that such evidence need not always be undermining: we are not wholly irrational creatures, after all. My independence condition provides a principled way of distinguishing between the innocuous and problematic sorts of irrelevant belief influence. It picks out ordinary, non-skeptical cases in which we get evidence of our own error. It says that, in those cases, evidence of irrelevant influence is epistemically significant. It also shows how self-doubt can be rational, and thus how accommodating evidence of our own error is compatible with the epistemic lives we see ourselves living. We are fallible creatures, yes, but we are also capable and intelligent ones. We can recognize and correct for our own error so as to improve our imperfect, yet nevertheless robust, epistemic lives.⁴³

⁴² There is a more general anti-skeptical strategy in this spirit. The strategy is most commonly attributed to Wittgenstein (1969) and Wright (2004). My goal here is not so ambitious. It is just to distinguish worries about irrelevant influence from skeptical worries. My proposal remains silent on how to address the latter.

⁴³ This paper benefited from many comments and conversations. I would especially like to thank Jennifer Acker, David J. Barnett, David Christensen, Tyler Doggett, Tom Dougherty, Adam Elga, Catherine Elgin, Jyl Gentzler, Alex George, Caspar Hare, Elizabeth Harman, Sophie Horowitz, Hilary Kornblith, Kathryn Lindeman, Amanda MacAskill, Ned Markosian, Joe Moore, Dilip Ninan, Alejandro Pérez Carballo, David Plunkett, Miriam Schoenfield, Nishi Shah, Paulina Sliwa, Justin Snedegar, Margot Strohminger, Sarah Stroud, Agustín Rayo, Judith Jarvis Thomson, Nick Treanor, Michael Veber, Kenny Walden, Dennis Whitcomb, Roger White and an anonymous referee. Thanks also to organizers and audiences at the New Insights Project: Etiology, Debunking, and Religious Epistemology Workshop, the Midsummer Philosophy Workshop, the 2012 Eastern APA, and the University of Vermont.

References

- Anur, Y. and D. Scott-Kakures, (forthcoming) “How Irrelevant Influences Bias Belief” in *Philosophical Perspectives*.
- Christensen, D. (2014) “Conciliation, Uniqueness, and Rational Toxicity” in *Noûs*, doi: 10.1111/nous.12077.
- (2011) “Disagreement, Question-Begging and Epistemic Self-Criticism” in *Philosophers Imprint*, 11(6): 1–22.
- (2010) “Higher-Order Evidence” in *Philosophy and Phenomenological Research*, 81 (1): 185–215.
- (2009) “Disagreement as Evidence: The Epistemology of Controversy” in *Philosophy Compass*, 4(1): 1–12.
- (2007) “Epistemology of Disagreement: The Good News” in *Philosophical Review*, 116: 187–217.
- Cohen, G.A. (2000) *If You’re an Egalitarian, How Come You’re so Rich?*, Harvard University Press, Cambridge, Massachusetts.
- De Cruz, H. (ms) “Irrelevant Influences and Philosophical Practice: A Qualitative Study”.
- DiPaolo, J. and R. Simpson (2015) “Indoctrination Anxiety and the Etiology of Belief” in *Synthese*, doi: 10.1007/s1129-015-0919-6.
- Dworkin, R. (2011) *Justice for Hedgehogs*, Harvard University Press, Cambridge, Massachusetts.
- (1996) “Objectivity and Truth: You’d Better Believe It” in *Philosophy and Public Affairs*, 25(2): 87–139.
- Elga, A. (ms) “Lucky to be Rational”.
- (2007) “Reflection and Disagreement” in *Noûs*, 41(3): 478–502.
- Garber, D. (2009) *What Happens After Pascal’s Wager: Living Faith and Rational Belief*, The 2009 Aquinas Lecture, Marquette University Press, Milwaukee.
- (2007) “*Religio Philosophi*: Some thoughts on God, Reason, and Faith” in L. Antony (ed.), *Philosophers Without Gods: Meditations on Atheism and the Secular Life*, Oxford University Press, Oxford: 32–40.
- Joyce, R. (2006) *The Evolution of Morality*, MIT Press, Cambridge, Massachusetts.
- Kelly, T. (2013) “Evidence Can Be Permissive” in M. Steup, J. Turri and E. Sosa (eds.), *Contemporary Debates in Epistemology*, Wiley-Blackwell, Oxford: 298–311.
- (2010) “Peer Disagreement and Higher Order Evidence” in R. Feldman and T. Warfield (eds.), *Disagreement*, Oxford University Press: Oxford: 111–174.
- (2005) “The Epistemic Significance of Disagreement” in J. Hawthorne and T. Gendler Szabo (eds.), *Oxford Studies in Epistemology, volume 1*, Oxford University Press, Oxford: 167–196.
- Leiter, B. (2004) “The Hermeneutics of Suspicion: Recovering Marx, Nietzsche, and Freud” in B. Leiter (ed.), *The Future for Philosophy*, Oxford University Press, Oxford: 74–105.
- Lewis, D. (1971) “Immodest Inductive Methods” in *Philosophy of Science*, 38(1): 54–63.
- Meacham, C. J. G. (2014) “Impermissive Bayesianism” in *Erkenntnis*, 79(6): 1185–1217.
- Mogensen, A. (2015) “Contingency Anxiety and the Epistemology of Disagreement” in *Pacific Philosophical Quarterly*, doi: 10.1111/pap1.12099.
- Plantinga, A. (2000) *Warranted Christian Belief*, Oxford University Press, Oxford.

- Rosen, G. (2001) "Nominalism, Naturalism, Epistemic Relativism" in *Philosophical Perspectives*, 15: 69–91.
- Schechter, J. (ms) "Luck, Rationality, and Explanation: A Reply to Elga's 'Lucky to be Rational'".
- Schoenfield, M. (2012) "Permission to Believe: Why Permissivism is True and What It Tells Us About Irrelevant Influences on Belief" in *Noûs*, 48(2): 193–218.
- Sher, G. (2001) "But I Could Be Wrong" in *Social Philosophy and Policy*, 18(2): 64–78.
- Street, S. (2015) "Does Anything Really Matter or Did We Just Evolve to Think So?" in A. Byrne, J. Cohen, G. Rosen and S. Shiffrin (eds.), *The Norton Introduction to Philosophy*, W. W. Norton, New York: 685–695.
- (2006) "A Darwinian Dilemma For Realist Theories of Value" in *Philosophical Studies*, 127(1): 109–166.
- Vavova, K. (2015) "Evolutionary Debunking of Morality" in *Philosophy Compass*, 10 (2): 104–116.
- (2014a) "Debunking Evolutionary Debunking" in R. Shafer-Landau (ed.), *Oxford Studies in Metaethics*, volume 9, Oxford University Press, Oxford: 76–101.
- (2014b) "Confidence, Evidence, and Disagreement" in *Erkenntnis*, 79(1): 173–183.
- Weatherson, B. (ms) "Do Judgments Screen Evidence?".
- Wittgenstein, L. (1969) *On Certainty*, G.E.M. Anscombe and G.H. von Wright (eds.), Basil Blackwell, New York.
- White, R. (2010) "You just Believe that Because..." in *Philosophical Perspectives*, 24 (1): 573–615.
- (2005) "Epistemic Permissiveness" in *Philosophical Perspectives*, 19(1): 445–459.
- Wright, C. (2004) "Wittgensteinian Certainties" in D. McManus (ed.), *Wittgenstein and Scepticism*, Routledge, New York: 22–55.